

On generating the ring of matrix semi-invariants

Gábor Ivanyos*

Youming Qiao†

K. V. Subrahmanyam ‡

August 10, 2015

Abstract

For a field \mathbb{F} , let $R(n, m)$ be the ring of invariant polynomials for the action of $\mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ on tuples of matrices – $(A, C) \in \mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ sends $(B_1, \dots, B_m) \in M(n, \mathbb{F})^{\oplus m}$ to $(AB_1C^{-1}, \dots, AB_mC^{-1})$. In this paper we call $R(n, m)$ the *ring of matrix semi-invariants*. Let $\beta(R(n, m))$ be the smallest D s.t. matrix semi-invariants of degree $\leq D$ generate $R(n, m)$. Guided by the Procesi-Razmyslov-Formanek approach of proving a strong degree bound for generating matrix invariants, we exhibit several interesting structural results for the ring of matrix semi-invariants $R(n, m)$ over fields of characteristic 0. Using these results, we prove that $\beta(R(n, m)) = \Omega(n^{3/2})$, and $\beta(R(2, m)) \leq 4$.

1 Introduction

Let \mathbb{F} be a field. In this article, we study the ring of invariant polynomials for the action of $\mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ on tuples of matrices – $(A, C) \in \mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$ sends $(B_1, \dots, B_m) \in M(n, \mathbb{F})^{\oplus m}$ to $(AB_1C^{-1}, \dots, AB_mC^{-1})$. Denoted by $R(n, m)$, we call this ring the *ring of matrix semi-invariants*, as (1) it is closely related to the classical ring of matrix invariants [Pro76] (see below for the definition, and [Dom00b, ANS07] for the precise relationship between these two rings); and (2) it is the ring of semi-invariants of the representation of the m -Kronecker quiver with dimension vector (n, n) . Here, the m -Kronecker quiver is the quiver with two vertices s and t , with m arrows pointing from s to t . When $m = 2$, it is the classical Kronecker quiver. The reader is referred to [DW00, SVdB01, DZ01] for a description of the semi-invariants for arbitrary quivers.

Let $\beta(R(n, m))$ be the smallest integer D s.t. matrix semi-invariants of degree $\leq D$ generate $R(n, m)$, and let $\sigma(R(n, m))$ be the smallest integer D s.t. matrix semi-invariants of degree $\leq D$ define the nullcone of $R(n, m)$.

Recently, the quantity $\sigma(R(n, m))$ has found several applications in computational complexity theory. In particular, if $\sigma(R(n, m))$ is polynomial in n and m , then it follows that (1) division gates can be efficiently removed in non-commutative arithmetic circuits with division ([HW14]); and (2) there exists a deterministic polynomial-time algorithm that decides whether the non-commutative rank of a square matrix of linear forms is full or not over the rational number field ([IQS15, Gur04]).

*Institute for Computer Science and Control, Hungarian Academy of Sciences, Budapest, Hungary. (Gabor.Ivanyos@szttaki.mta.hu).

†Centre for Quantum Computation and Intelligent Systems, University of Technology, Sydney, Australia. (jimmyqiao86@gmail.com)

‡Chennai Mathematical Institute, Chennai, India. (kv@cmi.ac.in).

We refer the interested reader to the cited works for further explanation on these connections. At present, the best bound for $\sigma(R(n, m))$ is $n!/\lceil n/2 \rceil!$ over large enough fields [IQS15].

One natural way to upper bound $\sigma(R(n, m))$ is of course to upper bound $\beta(R(n, m))$. Over algebraically closed fields of characteristic 0, by Derksen's result [Der01], $\beta(R(n, m)) \leq \max(2, 3/8 \cdot n^4 \cdot \sigma(R(n, m))^2)$. Therefore, if $\sigma(R(n, m))$ is polynomial in n then $\beta(R(n, m))$ is polynomial in n as well.¹ Another compelling reason to examine $\beta(R(n, m))$ is because of the following strong upper bound on β for the ring of matrix invariants over fields of characteristic 0.

Consider the action of $A \in \text{SL}(n, \mathbb{F})$ on $(B_1, \dots, B_m) \in M(n, \mathbb{F})^{\oplus m}$ by sending it to $(AB_1A^{-1}, \dots, AB_mA^{-1})$. The invariant ring w.r.t. this action is denoted as $S(n, m)$, and elements in $S(n, m)$ are called *matrix invariants*. The structure of $S(n, m)$ is well-understood over fields of characteristic 0: the first fundamental theorem (FFT), the second fundamental theorem (SFT), and an n^2 upper bound for $\beta(S(n, m))$ have been established in 1970's by Procesi, Razmyslov, and Formanek [Pro76, Raz74, For86]. Note that when applied to $S(n, m)$ over characteristic 0, Derksen's bound yields $\beta(S(n, m)) \leq \max(2, 3/8 \cdot n^2 \cdot \sigma^2)$ and $\sigma(S(n, m)) = n^{O(n^2)}$, far from the n^2 bound as mentioned above.

On the other hand, for $R(n, m)$, as far as we are aware, the best bound for $\beta(R(n, m))$ is $O(n^4 \cdot (n!)^2)$ over algebraically closed fields of characteristic 0, by combining the abovementioned results of [Der01] and [IQS15].

Our goal in this paper is to prove a better bound on $\beta(R(n, m))$ by following the approach of Procesi, Razmyslov, and Formanek. Therefore, in the following we restrict ourselves to fields of characteristic 0. While we do not achieve this, we describe several structural results for $R(n, m)$, including the second fundamental theorem (Proposition 1). Though some of these structural results should be known to experts, we could not find them in the literature, so we provide full proof details. Furthermore, these results allow us to prove that $\beta(R(n, m))$ has a lower bound $\Omega(n^{3/2})$ (Proposition 11).

One technical result that we believe is new, is Proposition 12. Roughly speaking, there exists a linear basis of matrix semi-invariants, such that each polynomial in this basis can be associated with a bipartite graph. An upper bound of D on $\beta(R(n, m))$ would follow, if we can prove that when the degree is $> D$, modulo the linear relations (as described in the second fundamental theorem), every polynomial can be written as a linear combination of those basis elements whose associated graphs are disconnected (Proposition 9). Proposition 12 then states that when $D > n^2$, every matrix semi-invariant of degree D can be written as a linear combination of those basis elements whose associated graphs are disconnected *or non-simple* (e.g. with at least one multiple edges) modulo the linear relations.

As an immediate consequence of Proposition 12, we prove that $\beta(R(2, m)) \leq 4$ over fields of characteristic 0 in Theorem 14. While this bound is known from Domokos' explicit generating set for $\beta(R(2, m))$ [Dom00a], we think this suggests the validity of the approach of Razmyslov, Procesi and Formanek when applied to matrix semi-invariants. Furthermore, since in this approach we do not exhibit invariants explicitly, we believe this approach will generalize to larger n .

More previous works. Here we collect a few more previous works on matrix semi-invariants. To start with, since as mentioned, matrix semi-invariants are just semi-invariants of the representations of the m -Kronecker quiver, results on semi-invariants of quivers apply to matrix semi-invariants, e.g.

¹Over fields of characteristic 0, by a theorem of Weyl [Wey97], $\beta(R(n, m)) \leq \beta(R(n, n^2))$. Therefore if $\beta(R(n, m))$ is polynomial in n and m then $\beta(R(n, m))$ is just polynomial in n . See [Dom03] for more on this.

the first fundamental theorem [DW00, SVdB01, DZ01]. Let us give one description from [DZ01]: suppose $R(n, m) \subseteq \mathbb{F}[x_{i,j}^{(k)}]$ where $i, j \in [n]$, $k \in [m]$, and $x_{i,j}^{(k)}$ are independent variables. Let $X_k = (x_{i,j}^{(k)})$. Then for $A_1, \dots, A_m \in M(d, \mathbb{F})$, $\det(A_1 \otimes X_1 + \dots + A_m \otimes X_m)$ is a matrix semi-invariant, and every matrix semi-invariant is a linear combination of such polynomials. Therefore, (B_1, \dots, B_m) is in the nullcone, if and only if for all $d \in \mathbb{Z}^+$ and all $(A_1, \dots, A_m) \in M(d, \mathbb{F})^{\oplus m}$, $A_1 \otimes B_1 + \dots + A_m \otimes B_m$ is singular.

A description of the nullcone of $R(n, m)$ can be found in [BD06, ANS07]. For certain small m or n , explicit generating sets of $R(n, m)$ have been computed in e.g. [Dom00b, Dom00a, DD12]. In these cases, elements of degree $\leq n^2$ generate the ring.²

Several results for matrix invariants over fields of positive characteristics are known: FFT was established by Donkin in [Don92, Don93], an n^3 upper bound for σ can be derived from [CIW97, Proposition 9], and Domokos in [Dom02, DKZ02] proved an upper bound $O(n^7 m^n)$ on β .

Organization. In Section 2 we briefly review the Procesi-Razmyslov-Formanek approach, and give an overview of our results. Then we describe the second fundamental theorem (Section 3), the \mathbb{FS}_{dn} bimodule structure (Section 4), and the S_{dn} diagonal action (Section 5). Finally, in Section 6, we use these structural results to prove a lower bound on $\beta(R(n, m))$, and that $\beta(R(2, m)) \leq 4$.

2 An overview of the structural results

An outline of the Procesi-Razmyslov-Formanek result. To motivate the results to be presented, we first review the n^2 bound for $S(n, m)$ over characteristic-0 fields by Razmyslov [Raz74] and Procesi [Pro76], and further elaborated by Formanek [For86]. Recall that $S(n, m)$ is the invariant ring of $A \in \mathrm{SL}(n, \mathbb{Q})$ on $(B_1, \dots, B_m) \in M(n, \mathbb{Q})^{\oplus m}$ by sending it to $(AB_1A^{-1}, \dots, AB_mA^{-1})$. Our exposition follows the one by Formanek [For86], and requires certain basic facts about the group algebra of the symmetric group as described therein.

Let X_i , $i \in [m]$ be an $n \times n$ matrix of indeterminates. Then $\mathrm{Tr}(X_{i_1} \cdot X_{i_2} \cdot \dots \cdot X_{i_k})$, $k \in \mathbb{Z}^+$, $i_j \in [m]$ generate all matrix invariants. The n^2 upper bound implies that invariants of this form with $k \leq n^2$ generate the ring of matrix invariants already.

The proof for this upper bound starts with identifying multilinear matrix invariants in the form $\mathrm{Tr}(X_{i_1} \cdot X_{i_2} \cdot \dots \cdot X_{i_k})$, $i_j \in [m]$ pairwise distinct, with the permutations of a single cycle (i_1, i_2, \dots, i_k) . Such identification gives a correspondence of multilinear invariants of degree d with the group algebra \mathbb{FS}_d .

The second step is to use the second fundamental theorem of matrix invariants, which suggests that the linear relations are spanned by the two-sided ideals indexed by partitions of length $> n$. At this point, it is clear that b is a degree bound, if and only if, for any $d > b$, the reducible elements (those permutations whose cycle types are not of a single cycle of length d) together with the linear relations span the whole space. Then, by using the standard bilinear form for CS_d (setting the permutations as an orthonormal basis), this is equivalent to showing that the space J spanned by the irreducible elements (those permutations whose cycle types are a single cycle of length d), and the space $I_{\leq n}$ spanned by the two-sided ideals indexed by partitions of length $\leq n$, intersect trivially.

²We thank M. Domokos for pointing out this fact to us.

The third step is to observe that the linear map f defined by $\sigma \rightarrow \text{sgn}(\sigma)\sigma$ where $\sigma \in S_d$, sends the two-sided ideal corresponding to λ , to the two-sided ideal corresponding to $\tilde{\lambda}$, the conjugate of λ . At the same time, f preserves every 1-dimensional subspace in J , as $\text{sgn}(\sigma)$ for any permutation $\sigma \in J$ is the same. Thus when $d > n^2$, J and $I_{\leq n}$ have to intersect trivially: by contradiction, suppose a nonzero $v \in J \cap I_{\leq n}$. Then $f(v) = \pm v$, thus $v \in I_{\leq n} \cap f(I_{\leq n})$. However, $f(I_{\leq n})$ is spanned by the two-sided ideals of width $\leq n$. When $d > n^2$, there does not exist a partition with both length and width $\leq n$. The n^2 degree bound then follows.

Overview of the structural results of $R(n, m)$. To carry on the above strategy for $S(n, m)$, a first step is to give the multilinear invariants in $R(n, m)$ a combinatorial description, which we take from [ANS07]. Since in this case we have multilinear invariants only when n divides m , we focus on $R(n, dn)$ in the following. Briefly speaking, we can identify a natural spanning set of multilinear invariants with n -regular bipartite graphs with d left (resp. right) vertices.

Then it is necessary to obtain the second fundamental theorem (SFT) for $R(n, dn)$. While certainly known to experts, we could not find an explicit statement in the literature, so we prove it in Proposition 1. Then we need to describe the $\mathbb{F}S_{dn}$ bimodule structure of $R(n, dn)$ (Fact 4). This is possible since $R(n, dn)$ can be viewed as a tensor product of two column tabloid modules [Ful97, Chap. 7.4]. So b is a degree bound, if whenever $dn > b$, with the help of relations, those multilinear invariants indexed by connected bipartite graphs can be written as a sum of those ones indexed by disconnected bipartite graphs (Proposition 9). Furthermore, a natural bilinear form on $R(n, dn)$ can be defined by setting the standard elements in $R(n, dn)$ as forming an orthonormal basis (Fact 4 (3)). This bilinear form makes sense w.r.t. the bimodule structure as well due to James' submodule theorem (Fact 3). Up to this point, we successfully parallel everything in the $S(n, m)$ setting.

What is missing is an analogue of the linear map f - this prevents us from completing this strategy. Nevertheless, what we've develop allows us to prove that for b to be a degree bound, then b must be $\Omega(n^{3/2})$ (Proposition 11).

To make progress, we exploit in depth the diagonal action of S_{dn} on $R(n, dn)$. Fortunately, the orbits under this action can be identified with row tabloid modules [Ful97, Chap. 7.2]. From this structure we obtain new forms of relations (Equation 3). We then prove that when $b > dn$, using these relations, each multilinear invariant indexed by a connected and simple (e.g. with no multiple edges) bipartite graph, can be written as a sum of those indexed by non-simple graphs (Proposition 12). This provides a non-trivial reduction result in the spirit of the degree bound statement.

We believe that the results we prove in this paper, will be useful to finally get a good degree bound for $R(n, m)$, over fields of characteristic 0.

3 The second fundamental theorem

Reduction to the multilinear case. Over characteristic 0 fields, the well-known two procedures, polarization and restitution, reduce many questions for general invariants to multilinear invariants.

For completeness, let us demonstrate this in the case of $R(n, m) \subseteq \mathbb{F}[x_{i,j}^{(k)}]_{k \in [m], i,j \in [n]}$. Let $X_k = (x_{i,j}^{(k)})_{i,j \in [n]}$. Recall that each $f \in R(n, m)$ has degree divisible by n . Suppose $\ell(X_1, \dots, X_m) = \sum_i a_i f_i$ forms a relation, where $f_i \in R(n, m)$. W.l.o.g. we can assume the f_i 's are multi-homogeneous, that is for every fixed $j \in [m]$, the degrees of f_i 's w.r.t. X_j are the same for every i .

Otherwise, we can divide ℓ into the multi-homogeneous components, and each component will again be a relation. Suppose $\deg(X_k)$ in ℓ is d_k , then $\sum_k d_k = dn$ for some $d \in \mathbb{Z}^+$. Then for each $k \in [m]$, we introduce d_k copies of X_k as $X_{(k,1)}, \dots, X_{(k,d_k)}$, as well as d_k variables $y_{k,1}, \dots, y_{k,d_k}$. Now consider the coefficient of $\prod_{i,j} y_{i,j}$ in $\ell(X_{(1,1)}y_{1,1} + \dots + X_{(1,d_1)}y_{1,d_1}, \dots, X_{(m,1)}y_{m,1} + \dots + X_{(m,d_m)}y_{m,d_m})$, and let it be ℓ' – the polarization of ℓ . It can be seen that ℓ' is a multilinear invariant in dn matrices, and ℓ' is a relation as well. Now we modify ℓ' by substituting $X_{(i,j)}$ by X_i , and let the result be ℓ'' – the substitution of ℓ' . Over characteristic 0 fields, we see that $\ell'' = \prod_i (d_i!) \ell$. As demonstrated above, since every relation can be obtained by restituting some multilinear relation, it is enough to understand multilinear relations.

Multilinear invariants of $R(n, m)$. As invariants in $R(n, m)$ are of degree divisible by n , $R(n, m)$ has a multilinear invariant if and only if n divides m . In the following we consider $R(n, m)$ for $m = dn$ where $d \in \mathbb{Z}^+$. We use the descriptions of invariants in $R(n, m)$ given in [ANS07]. We reformulate their results here.

We define a set of symbols $P = |i_1, \dots, i_n|$, $i_j \in [dn]$ with the anti-symmetric property:

$$|i_1, \dots, i_a, \dots, i_b, \dots, i_n| = -|i_1, \dots, i_b, \dots, i_a, \dots, i_n|.$$

Note that as \mathbb{F} is of characteristic 0, the anti-symmetric property implies that if there exist $i_j = i_k$ for $j \neq k$ then $|i_1, \dots, i_n| = 0$.

Then let \mathcal{P} be $\mathbb{F}\langle P \rangle$, the noncommutative polynomial ring with variables from P . A degree- d monomial in \mathcal{P} is of the form

$$|i_{1,1}, \dots, i_{1,n}| \circ \dots \circ |i_{d,1}, \dots, i_{d,n}|,$$

where \circ denotes the noncommutative product. We often record it as an $d \times n$ tableau

$$\begin{vmatrix} i_{1,1} & i_{1,2} & \cdots & i_{1,n} \\ i_{2,1} & i_{2,2} & \cdots & i_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ i_{d,1} & i_{d,2} & \cdots & i_{d,n} \end{vmatrix}. \quad (1)$$

A degree- d monomial $|i_{1,1}, \dots, i_{1,n}| \circ \dots \circ |i_{d,1}, \dots, i_{d,n}|$ is repetition-free, if $\{i_{j,k}, j \in [d], k \in [n]\} = [nd]$. Let $\mathcal{P}(d)$ be the \mathbb{F} -vector space of spanned by degree- d repetition-free monomials in \mathcal{P} .

Likewise we define another set of symbols $\widehat{P} = |\widehat{i}_1, \dots, \widehat{i}_n|$, $i_j \in [nd]$ satisfying also the anti-symmetric property, and define $\widehat{\mathcal{P}}$, and $\widehat{\mathcal{P}}(d)$ as before. The vector space $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ then has a basis $S \otimes \widehat{T}$, $S \in \mathcal{P}(d)$ and $\widehat{T} \in \widehat{\mathcal{P}}(d)$.

Now consider the vector space $(V \otimes W)^{\oplus dn}$ where $V \cong W \cong \mathbb{F}^n$. This is spanned by vectors of the form $(v_1 \otimes w_1) \oplus \dots \oplus (v_{dn} \otimes w_{dn})$, $v_i \in V$, $w_j \in W$. Given $i_1, i_2, \dots, i_n \in [dn]$, we define a function $\|i_1, \dots, i_n\|$, on $(v_1 \otimes w_1) \oplus \dots \oplus (v_{dn} \otimes w_{dn})$, as follows. $\|i_1, \dots, i_n\|$ sends $(v_1 \otimes w_1) \oplus \dots \oplus (v_{dn} \otimes w_{dn})$ to the determinant of the $n \times n$ matrix $[v_{i_1}, \dots, v_{i_n}]$. Likewise, define $\|\widehat{j}_1, \dots, \widehat{j}_n\|$, $j_k \in [dn]$ sending $(v_1 \otimes w_1) \oplus \dots \oplus (v_{dn} \otimes w_{dn})$ to the determinant of the matrix $[w_{j_1}, \dots, w_{j_n}]$. $\|i_1, \dots, i_n\|$ is known as a *Plücker coordinate*. We then consider the product function

$$\|i_{1,1}, \dots, i_{1,n}\| \cdots \|i_{d,1}, \dots, i_{d,n}\| \cdot \|\widehat{j}_{1,1}, \dots, \widehat{j}_{1,n}\| \cdots \|\widehat{j}_{d,1}, \dots, \widehat{j}_{d,n}\|$$

where $\{i_{k,\ell}, k \in [d], \ell \in [n]\} = \{j_{k,\ell}, k \in [d], \ell \in [n]\} = [nd]$, and extend this function to all of $(V \otimes W)^{\oplus dn}$ by linearity. This gives rise to a multilinear function on $(V \otimes W)^{\oplus dn}$.

We then define a linear map ϕ from $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ to multilinear functions on $(V \otimes W)^{\oplus dn}$, by sending

$$\begin{aligned} & (|i_{1,1}, \dots, i_{1,n}| \circ \dots \circ |i_{d,1}, \dots, i_{d,n}|) \otimes (|\widehat{j_{1,1}}, \dots, \widehat{j_{1,n}}| \circ \dots \circ |\widehat{j_{d,1}}, \dots, \widehat{j_{d,n}}|) \\ & \text{to } \|i_{1,1}, \dots, i_{1,n}\| \cdot \dots \cdot \|i_{d,1}, \dots, i_{d,n}\| \cdot \|\widehat{j_{1,1}}, \dots, \widehat{j_{1,n}}\| \cdot \dots \cdot \|\widehat{j_{d,1}}, \dots, \widehat{j_{d,n}}\| \end{aligned} \quad (2)$$

and extend by linearity. It is understood that if we write ϕ to apply to monomials in $\mathcal{P}(d)$ or $\widehat{\mathcal{P}}(d)$, we replace $|\cdot|$ to $\|\cdot\|$. It is not difficult to observe (see e.g. [ANS07]) that the image of ϕ are precisely the multilinear functions on $(V \otimes W)^{\oplus dn}$ which are invariant under the natural action $\text{SL}(n, \mathbb{F}) \times \text{SL}(n, \mathbb{F})$.

The formalism above is necessary for the formulation of the second fundamental theorem, which amounts to describe the kernel of ϕ .

Second fundamental theorem for multilinear invariants. In this part we describe the kernel of ϕ . It is clear that when $d = 1$, $\ker(\phi)$ is 0. In the following we assume $d \geq 2$.

Recall the Plücker relations for Plücker coordinates: for $i_1, \dots, i_n, j_1, \dots, j_n \in [nd]$ and $k \in [n]$, we have

$$\begin{aligned} & \|i_1, \dots, i_n\| \cdot \|j_1, \dots, j_n\| \\ &= \sum_{1 \leq s_1 < s_2 < \dots < s_k \leq n} \|i_1, \dots, j_1, \dots, j_k, \dots, i_n\| \cdot \|i_{s_1}, \dots, i_{s_k}, j_{k+1}, \dots, j_n\|, \end{aligned}$$

where j_1, \dots, j_k are in positions s_1, \dots, s_k . For a monomial T in $\mathcal{P}(d)$ recorded as in Equation 1, following Fulton [Ful97, pp. 97]³, we use $\pi_{j,k}(T)$ to denote the vector in $\mathcal{P}(d)$ of the form $\sum T'$, where T' runs over the tableaux obtained by switching the first k elements in the $j+1$ th row, with k elements in the j th row of T . Let $\mathcal{K}(d)$ be the span of all $T - \pi_{j,k}(T)$, where $T \in \mathcal{P}(d)$, $j \in [n-1]$, and $k \in [n]$. In particular, note that when $k = n$, $\pi_{j,n}(T)$ is just switching the j th and $j+1$ th row. So though monomials in $\mathcal{P}(d)$ are non-commutative, modulo $\mathcal{K}(d)$ they become commutative, which is consistent with the image of ϕ .

A monomial T in $\mathcal{P}(d)$ is *standard*, if its tableau is (strictly) increasing in each row and each column. Following a procedure called the straightening, it is well-known that the standard monomials form a basis for the quotient space $\mathcal{P}(d)/\mathcal{K}(d)$.

Similarly we have $\widehat{\mathcal{K}}(d)$ in $\widehat{\mathcal{P}}(d)$. By Plücker relations, we know that $\langle \mathcal{K}(d) \otimes \widehat{\mathcal{P}}(d) \cup \mathcal{P}(d) \otimes \widehat{\mathcal{K}}(d) \rangle$ lies in $\ker(\phi)$. (Recall that $\langle \cdot \rangle$ denotes the linear span.) We show that these two sets are equal.

Proposition 1. *Let notations be as above. We have $\langle \mathcal{K}(d) \otimes \widehat{\mathcal{P}}(d) \cup \mathcal{P}(d) \otimes \widehat{\mathcal{K}}(d) \rangle = \ker(\phi)$.*

Proof. It remains to prove that $\ker(\phi) \subseteq \langle \mathcal{K}(d) \otimes \widehat{\mathcal{P}}(d) \cup \mathcal{P}(d) \otimes \widehat{\mathcal{K}}(d) \rangle$. Let ℓ be in $\ker(\phi)$. Suppose ℓ is

$$c_1 S_1 \otimes \widehat{T}_1 + c_2 S_2 \otimes \widehat{T}_2 + \dots + c_k S_k \otimes \widehat{T}_k$$

where c_i 's are in \mathbb{F} , and S_i 's (resp. \widehat{T}_i 's) are monomials in $\mathcal{P}(d)$ (resp. $\widehat{\mathcal{P}}(d)$). Note that S_i 's (resp. \widehat{T}_i 's) are not necessarily distinct. We first arrange according to S_i 's, and write ℓ as

$$S_1 \otimes h_1 + S_2 \otimes h_2 + \dots + S_{k'} \otimes h_{k'}$$

³With a slight change from column as in [Ful97] to row here.

where $h_i \in \widehat{\mathcal{P}}(d)$, and $S_1, \dots, S_{k'}$ are distinct. Modulo the space $\mathcal{K}(d) \otimes \widehat{\mathcal{P}}(d)$, we express each S_i as a linear combination of those ones indexed by standard tableaux. Then by re-grouping according to the standard tableaux, ℓ is expressed as

$$S'_1 \otimes h'_1 + S'_2 \otimes h'_2 + \dots + S'_p \otimes h'_p$$

where $h'_i \in \widehat{\mathcal{P}}(d)$, and S'_1, \dots, S'_p are standard and distinct. Now we claim that every $S'_i \otimes h'_i$ must be in $\mathcal{P}(d) \otimes \widehat{\mathcal{K}}(d)$. W.l.o.g. assume $S'_1 \otimes h'_1$ is not, that is h'_1 is not in $\widehat{\mathcal{K}}(d)$. Then $\phi(h'_1)$ is not a zero function on $W^{\oplus nd}$. So there exists some $\mathbf{w} = w_1 \oplus w_2 \oplus \dots \oplus w_{nd}$ s.t. $\phi(h'_1)(\mathbf{w}) \neq 0$. Now we restrict $\phi(\ell)$ to the set $(V \otimes w_1) \oplus \dots \oplus (V \otimes w_{nd})$ which yields

$$\phi(S'_1) \cdot \phi(h'_1)(\mathbf{w}) + \phi(S'_2) \cdot \phi(h'_2)(\mathbf{w}) + \dots + \phi(S'_j) \cdot \phi(h'_j)(\mathbf{w}).$$

Since S'_i 's are standard and distinct, $\phi(S'_i)$'s are linearly independent as functions on $V^{\oplus dn}$. As $\phi(h'_1)(\mathbf{w}) \neq 0$, the restriction of ℓ on $(V \otimes w_1) \oplus \dots \oplus (V \otimes w_{nd})$, as a function on $V^{\oplus d}$, is nonzero. This contradicts ℓ being a zero function on $(V \otimes W)^{\oplus dn}$. So every $S'_i \otimes h'_i$ is in $\mathcal{P} \otimes \widehat{\mathcal{K}}$. From this we conclude that $\ell \in \langle \mathcal{K}(d) \otimes \widehat{\mathcal{P}}(d) \cup \mathcal{P}(d) \otimes \widehat{\mathcal{K}}(d) \rangle$. \square

4 The \mathbb{FS}_{dn} bimodule structure

There exists a natural action of $S_{dn} \times S_{dn}$ on $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$. Suppose S and \widehat{T} are monomials in \mathcal{P} and $\widehat{\mathcal{P}}$ respectively. Then $(\sigma, \tau) \in S_{dn} \times S_{dn}$ sends $S \otimes \widehat{T}$ to $S^\sigma \otimes \widehat{T}^\tau$, where σ and τ permute the entries of S and \widehat{T} . This endows $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ an \mathbb{FS}_{dn} bimodule structure. This action descends to the relations describing the second fundamental theorem: if ℓ is a relation (in $\ker(\phi)$), then $\ell^{(\sigma, \tau)}$ is also a relation. So $\ker(\phi)$ is an \mathbb{FS}_{dn} sub-bimodule, whose structure will be described in the following Fact 4.

Such an action on S_{dn} on $\mathcal{P}(d)$ is well-understood: this is the column tabloid module as discussed in [Ful97, Chap. 7.4].⁴ The column tabloid module is a dual of the more well-known row tabloid module as shown in [Ful97, Chap. 7.2]. We collect some basic facts about the column tabloid module adapted to our setting. Recall that for a nonnegative integer s , a partition of size s is a non-increasing sequence of nonnegative integers $\lambda = (\lambda_1, \dots, \lambda_\ell)$, with $\sum_{i=1}^\ell \lambda_i = s$. This is denoted as $\lambda \vdash s$. We identify $(\lambda_1, \dots, \lambda_\ell)$ with $(\lambda_1, \dots, \lambda_\ell, 0, \dots, 0)$; namely any trailing zeros are assumed implicitly if required. The conjugate of λ denoted by $\tilde{\lambda}$ is a partition of s with $\tilde{\lambda}_i = |\{j \in [\ell] \text{ s.t. } \lambda_j \geq i\}|$. The height of λ is $\max\{i \in [\ell] \mid \lambda_i \neq 0\}$, and the width is λ_1 . Partitions are usually represented using Young diagrams: that is a concatenation of rows of boxes arranged to be left aligned, with the i th row having λ_i boxes. For two partitions $\nu = (\nu_1, \dots, \nu_k)$ and $\lambda = (\lambda_1, \dots, \lambda_\ell)$, ν dominates λ if for any $j \in \mathbb{Z}^+$, $\sum_{i=1}^j \nu_i \geq \sum_{i=1}^j \lambda_i$. We use d^n to denote the partition (d, \dots, d) of height n . For a partition λ of $[dn]$, $\text{Sp}(\lambda)$ denotes the Specht module corresponding to λ over \mathbb{F} .

Fact 2 ([Ful97, Chap. 7.4]). 1. As an \mathbb{FS}_{dn} module, $\mathcal{P}(d)$ decomposes as

$$\text{Sp}(d^n) \oplus \left(\bigoplus_{\lambda \vdash dn} \text{Sp}(\lambda)^{\oplus c_\lambda} \right),$$

⁴Note that the notation in [Ful97] is a bit different from ours; there the column tableaux satisfy anti-symmetric property along the columns. So one needs to flip the tableaux here to match the results there.

where λ runs over partitions of dn strictly dominated by d^n , and $c_\lambda = K_{\tilde{\lambda}, n^d}$, the Kostka number w.r.t. $\tilde{\lambda}, n^d$.

2. Given a monomial $S \in \mathcal{P}(d)$ represented by a tableau, the column subgroup of S , $\text{cl}(S)$ is the subgroup of S_{dn} that preserves the columns of S . The (unsigned) column symmetrizer $c(S) = \sum_{\pi \in \text{cl}(S)} \pi$. Then $\mathcal{E}(d) := \langle c(S) \cdot S \mid S \in \mathcal{P}(d) \text{ are monomials} \rangle \cong \text{Sp}(d^n)$.
3. $\mathcal{K}(d)$ is a submodule of $\mathcal{P}(d)$ and is isomorphic to $\bigoplus_{\lambda \vdash dn} \text{Sp}(\lambda)^{\oplus c_\lambda}$.

We also observe the following adaptation of James' submodule theorem to column tabloid modules.

Fact 3 (See e.g. [Sag01, Theorem 2.4.4]). *Let β' be the bilinear form on $\mathcal{P}(d)$ by setting the monomials as an orthonormal basis. Then $\mathcal{K}(d)$ and $\mathcal{E}(d)$ are orthogonal complement of each other under β' .*

Thus, as an $\mathbb{F}S_{dn}$ bimodule, $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ is the tensor product of two column tabloid modules. Its structure is then easily deduced from Fact 2 and 3.

Fact 4. 1. As an $\mathbb{F}S_{dn}$ bimodule, $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ decomposes as

$$(\text{Sp}(d^n) \otimes \text{Sp}(d^n)) \oplus \left(\bigoplus_{\lambda, \nu \vdash dn} (\text{Sp}(\lambda) \otimes \text{Sp}(\nu))^{\oplus c_\lambda \cdot c_\nu} \right),$$

where λ and ν are dominated by d^n , and at least one of λ and ν is strictly dominated by d^n . c_λ and c_ν are Kostka numbers as in Fact 2.

2. $\ker(\phi) = \langle \mathcal{K}(d) \otimes \widehat{\mathcal{P}}(d) \cup \mathcal{P}(d) \otimes \widehat{\mathcal{K}}(d) \rangle$ is an $\mathbb{F}S_{dn}$ sub-bimodule, isomorphic to the summand $\bigoplus_{\lambda, \nu \vdash dn} (\text{Sp}(\lambda) \otimes \text{Sp}(\nu))^{\oplus c_\lambda \cdot c_\nu}$, as above.
3. Let β be the bilinear form on $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ by setting $S \otimes \widehat{T}$, where S and \widehat{T} are monomials in $\mathcal{P}(d)$ and $\widehat{\mathcal{P}}(d)$, respectively. Then $\mathcal{E}(d) \otimes \widehat{\mathcal{E}}(d)$ and $\langle \mathcal{P}(d) \otimes \widehat{\mathcal{K}}(d) \cup \mathcal{K}(d) \otimes \widehat{\mathcal{P}}(d) \rangle$ are orthogonal complement of each other under β .

5 The diagonal action of S_{dn}

In this subsection we consider the diagonal action of S_{dn} on $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$. That is, $\sigma \in S_{dn}$ acts on $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ as $(\sigma, \sigma) \in S_{dn} \times S_{dn}$.

To understand this structure we introduce a combinatorial structure called the *correlated tableaux*. A correlated tableau C of size $d \times d$ is a $d \times d$ matrix whose entries are subsets of $[dn]$, satisfying: (1) $C(i, j)$, $i, j \in [d]$ form a partition of $[dn]$; (2) $\forall i \in [d]$, $|\cup_{j \in [d]} C(i, j)| = n$; (3) $\forall j \in [d]$, $|\cup_{i \in [d]} C(i, j)| = n$.

Correlated tableaux are in 1-to-1 correspondence with $S \otimes \widehat{T}$ where S and \widehat{T} are monomials from $\mathcal{P}(d)$ and $\widehat{\mathcal{P}}(d)$, respectively. Given $S \otimes \widehat{T}$, we can form a correlated tableau C by setting $C(i, j) = \{k \in [nd] \mid k \in S_i \text{ and } \hat{k} \in T_j\}$. Given a correlated tableau C , by reading along the rows from top to bottom we get a monomial S , and by reading along the columns from left to right (and adding $\hat{\cdot}$) we get a monomial \widehat{T} .

Let us set up the following convention about arranging the rows of S (resp. \widehat{T}) as in $S \otimes \widehat{T}$, where S and \widehat{T} are understood as tableaux as shown in Equation 1. In the literature, it is more common to arrange each row of S to be increasing, namely $i_{j,1} < i_{j,2} < \dots < i_{j,n}$ for every j . But when S appears in $S \otimes \widehat{T}$ we shall order each row while taking into consideration of \widehat{T} as well. That is, first form the correlated tableau C based on $S \otimes \widehat{T}$, and then read each row of C from left to right. If there are more than 1 elements in an entry, then use the increasing order. The rows of \widehat{T} are also ordered in a similar way.

As an example, consider

$$S = \begin{vmatrix} 1 & 2 & 4 \\ 3 & 5 & 6 \\ 7 & 8 & 9 \end{vmatrix}, \widehat{T} = \begin{vmatrix} \widehat{1} & \widehat{2} & \widehat{7} \\ \widehat{3} & \widehat{4} & \widehat{5} \\ \widehat{6} & \widehat{8} & \widehat{9} \end{vmatrix}.$$

The correlated tableaux then is $C =$

{1, 2}	{4}	\emptyset
\emptyset	{3, 5}	{6}
{7}	\emptyset	{8, 9}

. So we obtain by reading C

$$S' = \begin{vmatrix} 1 & 2 & 4 \\ 3 & 5 & 6 \\ 7 & 8 & 9 \end{vmatrix}, \widehat{T}' = \begin{vmatrix} \widehat{1} & \widehat{2} & \widehat{7} \\ \widehat{4} & \widehat{3} & \widehat{5} \\ \widehat{6} & \widehat{8} & \widehat{9} \end{vmatrix}.$$

That is $\widehat{T}' = -\widehat{T}$.

Having defined correlated tableaux, we define the (d, n) -correlated diagram as a $d \times d$ matrix D over \mathbb{N} with row and column sums being n . For a correlated diagram D , we form its *diagonal partition* $\delta(D)$, by arranging the entries of D to be non-increasing to get a partition of dn . Each correlated tableau C gives rise to a correlated diagram D_C by setting $D_C(i, j) = |C(i, j)|$; we all call

D_C the shape of C . Using the example as above, $D_C =$

2	1	0
0	2	1
1	0	2

. $\delta(D_C)$ then is $(2, 2, 2, 1, 1, 1)$.

Fix a correlated diagram D , and let $\mathcal{C}(D)$ be the subspace of $\mathcal{P}(d)$, spanned by all correlated tableaux of shape D . $\mathcal{C}(D)$ is clearly a submodule of $\mathcal{P}(d)$ under the diagonal action of S_{dn} . Let us define the *row tabloid module* of $\mathbb{F}S_s$, $s \in \mathbb{Z}^+$. For a partition $\lambda \vdash s$, row tabloids of shape λ are fillings of the Young diagram of shape λ by $[s]$ without repetitions, with the following equivalence relation: two row tabloids P and Q are equivalent if the corresponding rows have the same entries. Row tabloids admit an $\mathbb{F}S_s$ action by permuting the entries. The linear span of all row tabloids of shape λ is then an $\mathbb{F}S_s$ module and denoted as $\mathcal{R}(\lambda)$. The decomposition of the row tabloid module is well-known and quite similar to that of the column tabloid module as shown in Fact 2; see [Ful97, Chap. 7.2].

Lemma 5. *As an $\mathbb{F}S_{dn}$ module, $\mathcal{C}(D)$ is isomorphic to the row tabloid module $\mathcal{R}(\delta(D))$.*

Proof. To set up a linear map between $\mathcal{C}(D)$ and $\mathcal{R}(\delta(D))$, arrange the entries of $C \in \mathcal{C}(D)$ following an arbitrary but fixed order as long as this order maintains the sizes of the entries to be non-increasing. This gives a row tabloid in $\mathcal{R}(\delta(D))$, and then extend by linearity. To see that the actions of S_{dn} are compatible, note that by our convention of sending correlated tableaux to monomials, the diagonal action of S_{dn} is just permuting the entries in the correlated tableau. In particular, if i and j are in the same entry of C , then switching i and j leaves C unchanged, as two -1 are produced from the two monomials associated with C . \square

Proposition 6. *As an \mathbb{FS}_{dn} module, $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ is isomorphic to $\oplus_D \mathcal{R}(\delta(D))$, where D runs over all (d, n) -correlated diagrams.*

We also note that the well-known Robinson-Schensted-Knuth correspondence gives a 1-to-1 correspondence between (d, n) -correlated diagrams and pairs of semistandard tableaux with the same shape λ , and the same content as n^d . Note that λ necessarily dominates n^d to satisfy the semistandard condition. On the other hand, such pairs of semistandard tableaux could be used to index those \mathbb{FS}_{dn} sub-bimodules of $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ of isomorphic type $\mathrm{Sp}(\tilde{\lambda}) \otimes \mathrm{Sp}(\tilde{\lambda})$. This follows by adapting the results in [Sag01, Chap. 2.10] to the column tabloid setting.

Recall that ϕ realizes $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ as functions on $(V \otimes W)^{\oplus dn}$, as defined in Equation 2. We now describe another set of vectors in $\ker(\phi)$ based on the \mathbb{FS}_{dn} structure. Given a correlated tableau C , suppose C has $\geq n^2 + 1$ nonempty entries. Note that this requires $d \geq n + 1$. Fix $N := n^2 + 1$ entries and from each entry choose a number, denoted as $I = \{i_1, \dots, i_{n^2+1}\}$. Then consider the following vector in $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$:

$$\mathrm{Alt}(C, I) = \sum_{\sigma \in S_I} \mathrm{sgn}(\sigma) C^\sigma, \quad (3)$$

where σ acts on C as in the diagonal action of S_{dn} . Then $\phi(\mathrm{Alt}(C, I))$ is a zero function on $(V \otimes W)^{\oplus dn}$, since it is alternating in $n^2 + 1$ copies of $V \otimes W$, an n^2 -dimensional vector space. Note that construction is very natural in light of the identification of $\mathcal{C}(D)$ with $\mathcal{R}(\delta(D))$.

6 Towards proving degree bounds

Reduction to the multilinear case. Over characteristic 0 fields, we can also reduce the degree bound problem to the multilinear case as well. Consider a graded invariant ring R , for which we want to prove that R is generated as a ring by $R_{\leq b} = \{f \in R \mid \deg(f) \leq b\}$. If R is generated by $R_{\leq b}$, then any homogeneous multilinear invariant f of degree $b' > b$ is equal to $\sum_i \prod_j f_{i,j}$ where $b \geq \deg(f_{i,j}) \geq 1$. On the other hand, suppose each multilinear invariant of degree $b' > b$ can be written as such. Take any homogeneous $g \in R$ of degree $b' > b$, fully polarize it to get a multilinear f , which then by assumption can be written as $\sum_i \prod_j f_{i,j}$ where $b \geq \deg(f_{i,j}) \geq 1$. Then reconstitute back to get $b'!g$ on one hand, and on the other hand $\sum_i \prod_j g_{i,j}$, where $g_{i,j}$ are obtained via restitution to $f_{i,j}$. So $b \geq \deg(g_{i,j}) = \deg(f_{i,j}) \geq 1$, which implies that R is generated by $R_{\leq b}$.

Some reducible polynomials in $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$. Recall that $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ is spanned by $S \otimes \widehat{T}$, where S and \widehat{T} are monomials; or equivalently, by correlated tableaux C . We now view C as a bipartite graph on $(L \cup R, E)$ allowing multiple edges, where $L = R = [d]$, and there are k edges between i and j if $C(i, j) = k$.

We shall prove that $\phi(C)$ is reducible if and only if this graph is disconnected. To show this, we recall the matrix-theoretic interpretation of $\phi(C)$ as in [ANS07]. Let A_1, \dots, A_s be s variable matrices of size $s \times s$, and y_1, \dots, y_s are newly introduced variables. The *mixed discriminant* mdisc is a polynomial in s^3 variables in A_i 's, defined as follows. Given $\sigma \in S_s$, let A^σ be the matrix whose i th column is the i th column of $A_{\sigma(i)}$. Then $\mathrm{mdisc}(A_1, \dots, A_s) = \sum_{\sigma \in S_s} \det(A^\sigma)$.

Given a correlated tableau C , we form dn matrices of size $dn \times dn$ Y_1, \dots, Y_{dn} as follows. For every i , Y_i is viewed a $d \times d$ block matrix, where each block is of size $n \times n$. If i appears in the

(j, k) th position of C , then Y_i has X_i at the (j, k) block, and 0 everywhere else. In particular note that by construction $\text{mdisc}(Y_1, \dots, Y_{dn})$ is a multilinear polynomial in X_i 's.

Theorem 7 ([ANS07]). $\phi(C)$ is equal to $\text{mdisc}(Y_1, \dots, Y_{dn})$.

Proposition 8. $\phi(C)$ is a reducible polynomial if and only if C is a disconnected graph.

Proof. If C is disconnected, then $\phi(C)$ is a product of at least two multilinear polynomials of smaller degree, each of which corresponds to a connected component. On the other hand, if C is connected, let us assume that $\phi(C)$ can be written as $f \cdot g$. Without loss of generality assume f is not constant, and let us show that g is constant. As f is not constant, some variable $X_i(a, b)$ appears in f . Suppose a variable $X_p(c, d)$ appearing in $\phi(C)$ does not appear with $X_i(a, b)$ in any monomial. Since $\phi(C)$ is multilinear and $X_j(c, d)$ appears in $\phi(C)$, the only way this can happen is if $X_j(c, d)$ is also in f . We say $X_j(c, d)$ connects to $X_i(a, b)$ if this happens. Note that if $X_k(u, v)$ connects to $X_j(c, d)$ and $X_j(c, d)$ connects to $X_i(a, b)$, then $X_k(u, v)$ is in f as well. It is not hard to see from mixed discriminant perspective that: (1) every variable in X_i connects to $X_i(a, b)$; (2) if X_j is in the same row or column as X_i in C , then every variable in X_j connects to some variable in X_i . Since C is a connected graph, using (1) and (2) iteratively we put every variable in f , which implies that g is a constant. \square

Let $\mathfrak{D}(d)$ be the span of disconnected correlated tableaux in $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$, and $\mathfrak{S}(d)$ the span of connected correlated tableaux.

Using the second fundamental theorem. A degree bound of, say, degree bn , would involve showing that every invariant of degree dn , $d > b$, is equal to a sum of products of invariants of smaller degree, modulo relations in ϕ . Algorithmically we would start with an invariant of degree d , and express it as a sum of products of invariants of smaller degree plus something in the kernel ϕ . If in this new expression there still remain invariants of degree bigger than b , we write those as sums of products of invariants of smaller degree plus elements in the kernel ϕ . We continue, till all invariants involved in the final expression have degree less than or equal to bn .

We are then led to examine the following situation – if B is a set of polynomials spanning $R(n, dn)$, is every irreducible polynomial in B equal to a sum of reducible ones in B plus something in the kernel of ϕ ? If this happens beyond a certain degree bn , we get an upper bound of bn , on the degree in which the invariant ring is generated. Taking B as $\phi(C)$, where C runs over all correlated tableaux, we summarise the above argument in the following

Proposition 9. $R(n, m)$ is generated by $R(n, m)_{\leq bn}$ if and only if for every $d > b$, $\mathfrak{D}(d) \cup \ker(\phi)$ spans $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$.

Using the bilinear form β introduced in Fact 4 (3), we have the following.

Corollary 10. $R(n, m)$ is generated by $R(n, m)_{\leq bn}$ if and only if for every $d > b$, $\mathfrak{C}(d) \cap (\mathcal{E}(d) \otimes \mathcal{E}(d)) = 0$.

We now use some results developed so far to lower bound the degree in which the invariant ring can be generated as an algebra.

Proposition 11. If $R(n, m)$ is generated in degree bn , then $b = \Omega(n^{1/2})$.

Proof. We use proposition 9. First observe that for $\mathfrak{D}(d) \cup \ker(\phi)$ to span $\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)$ it is necessary that $\dim(\mathfrak{D}(d)) + \dim(\ker(\phi)) \geq \dim(\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d))$. We show that when $d < o(n^{1/2})$ this does not happen, by doing a dimension count. To compute the dimensions proceed as follows.

1. $\dim(\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)) = D_1 := \left(\frac{(nd)!}{(n!)^d}\right)^2$;
2. By the hook length formula, $\dim(\mathcal{E}(d) \otimes \mathcal{E}(d)) = D_2 := \left(\frac{(nd)!}{(1 \cdot \dots \cdot n) \cdot (2 \cdot \dots \cdot (n+1)) \cdot \dots \cdot (d \cdot \dots \cdot (d+n-1))}\right)^2$;
3. $\dim(\ker(\phi)) = D_1 - D_2$;
4. A lower bound on $\dim(\mathfrak{D}(d))$ is $\left(\frac{(n(d-1))!}{(n!)^{d-1}}\right)^2$. This counts only those disconnected graphs with $C(1, 1) = \{1, \dots, n\}$.
5. An upper bound on $\dim(\mathfrak{D}(d))$ is

$$\sum_{k=1}^{d-1} \binom{d}{k}^2 \cdot \binom{dn}{kn} \cdot \left(\frac{(kn)!}{(n!)^k}\right)^2 \cdot \left(\frac{((d-k)n)!}{(n!)^{d-k}}\right)^2,$$

where $\binom{d}{k}^2$ fixes k rows and columns as a connected component C , $\binom{dn}{kn}$ chooses kn numbers as the labels within C , and the rest terms count the number of ways filling in C and the other component.

Using (3) and (5), we prove that when $d = o(\sqrt{n})$ and n is large enough, then $\dim(\mathfrak{D}(d)) < \dim(\mathcal{P}(d) \otimes \widehat{\mathcal{P}}(d)) - \dim(\ker(\phi))$, which implies that the degree bound is necessarily $\Omega(n^{3/2})$. For this we need to verify that

$$\sum_{k=1}^{d-1} \binom{d}{k}^2 \cdot \binom{dn}{kn} \cdot \left(\frac{(kn)!}{(n!)^k}\right)^2 \cdot \left(\frac{((d-k)n)!}{(n!)^{d-k}}\right)^2 < \left(\frac{(nd)!}{(1 \cdot \dots \cdot n) \cdot (2 \cdot \dots \cdot (n+1)) \cdot \dots \cdot (d \cdot \dots \cdot (d+n-1))}\right)^2$$

for $d = o(\sqrt{n})$. It is easy to see that this is equivalent to showing that

$$\sum_{k=1}^{d-1} \binom{d}{k}^2 \cdot \frac{(kn)!((d-k)n)!}{(dn)!} < \left(\frac{(n!)^d}{(1 \cdot \dots \cdot n) \cdot (2 \cdot \dots \cdot (n+1)) \cdot \dots \cdot (d \cdot \dots \cdot (d+n-1))}\right)^2.$$

Upper bounding the left hand side by $2^{2d} \cdot \frac{n!((d-1)n)!}{(dn)!}$ and rearranging the terms we need to show that

$$\frac{(dn)!}{n!((d-1)n)!} > \left(2^d \cdot \frac{(1 \cdot \dots \cdot n) \cdot (2 \cdot \dots \cdot (n+1)) \cdot \dots \cdot (d \cdot \dots \cdot (d+n-1))}{(n!)^d}\right)^2.$$

The left hand side is at least $\frac{(dn-n+1) \cdot \dots \cdot (dn)}{n^n}$ and the right hand side is at most

$$[2^d \cdot (n+1)^{d-1} \cdot (n+2)^{d-2} \cdot \dots \cdot (n+d-1)]^2,$$

which is in turn at most

$$[2^d \cdot (n+d-1) \binom{d}{2}]^2.$$

It is easy to verify now that when $d = o(\sqrt{n})$, $(dn - n + 1) \cdot \dots \cdot (dn)$ is asymptotically larger than $n^n \cdot [2^d \cdot (n+d-1) \binom{d}{2}]^2$. \square

On the other hand, using (4) and (5) it is easy to deduce that when $d = \omega(\sqrt{n})$ then $\dim(\mathfrak{D}(d)) > D_2$ asymptotically in n , and when $d > 4n$ then $\dim(\mathfrak{D}(d)) > D_2$ unconditionally. This means that once d is $\Omega(n)$, then $\dim(\mathfrak{D}(d)) + \dim(\ker(\phi)) \geq \dim(\mathcal{P}(d) \otimes \widehat{\mathcal{P}(d)})$, clearing a bottleneck to prove a polynomial degree bound for $R(n, m)$.

On connected graphs without multiple edges. Recall that by Proposition 9 we need to consider whether $\mathfrak{D}(d) \cup \ker(\phi)$ spans $\mathcal{P}(d) \otimes \widehat{\mathcal{P}(d)}$ for large enough d . In this section we prove an analogous, but weaker result.

Let $\mathfrak{S}(d)$ be the span of those correlated tableaux that are connected, and simple; that is, with no multiple edges. $\mathfrak{S}(d)$ is clearly a subspace of $\mathfrak{C}(d)$. Let $\mathfrak{M}(d)$ be the span of the tableaux that are not in $\mathfrak{S}(d)$. Namely those tableaux spanning $\mathfrak{M}(d)$ are either disconnected, or connected and with at least one multiple edge. In particular they include any tableau with at least one multiple edge.

In the proof of the following proposition, besides the relations $\mathcal{K}(d) \otimes \widehat{\mathcal{P}(d)}$ and $\mathcal{P}(d) \otimes \widehat{\mathcal{K}(d)}$, we crucially use those ones as described in Equation 3. We then refer to relations in $\mathcal{K}(d) \otimes \widehat{\mathcal{P}(d)}$ and $\mathcal{P}(d) \otimes \widehat{\mathcal{K}(d)}$ as the first type, and relations from Equation 3 as second type realtions.

Proposition 12. *Let $d \geq n + 1$. Then $\mathcal{P}(d) \otimes \widehat{\mathcal{P}(d)}$ is spanned by $\mathfrak{M}(d) \cup \ker(\phi)$.*

Note that if we define $\mathfrak{U}(d)$ the span of correlated tableaux whose graphs have at least one multiple edge, then this proposition is just to say that when $d > n$, $\mathfrak{D}(d) \cup \ker(\phi) \cup \mathfrak{U}(d)$ spans $\mathcal{P}(d) \otimes \widehat{\mathcal{P}(d)}$.

Proof. Take any correlated tableau $C \in \mathfrak{S}(d)$; namely the bipartite graph associated to C is connected and without multiple edges. We shall show that C can be written as a linear combination of tableaux with multiple edges via the help of $\ker(\phi)$.

To do that let G be the bipartite graph of C , and T be the set of edges of G . (T is of course labelled by $[nd]$, but we make T explicitly for clarity.) Note that $|T| = dn \geq n^2 + 1$. For some $\sigma \in S_T$, suppose $\sigma = \tau_1 \tau_2 \dots \tau_k$ where τ_i is a transposition switching e and e' in T . As G is connected, for each pair of edges $e, e' \in T$, e is connected to e' by some path in G . Thus τ_i can be further decomposed as a product of transpositions consisting of edges along that path connecting e and e' . That is for each i we have $\tau_i = \pi_{i,1} \pi_{i,2} \dots \pi_{i,j_i}$, where $\pi_{i,j}$ switches two edges sharing a common vertex.


Consider now $\pi_{1,1}$, switching adjacent edges e and e' , and suppose the position of e is (s, t_1) and e' is (s, t_2) , $s, t_1, t_2 \in [d]$ and $t_1 < t_2$. That is, we assume e and e' are in the same row as in the correlated tableau. By the relations of the first type, we move e to the t_2 th column along the s th row and get $C = -C^{\pi_{1,1}} + \sum_i D_i$, where $D_i \in \mathfrak{M}(n, d)$. We explain the terms on the RHS: the -1 sign before $C^{\pi_{1,1}}$ is because when we switch e and e' , the order of reading the r th row changed: from first reading e and then e' to first e' and then e . Note that $C^{\pi_{1,1}}$ is of the same shape as C . D_i 's are in $\mathfrak{M}(n, d)$ because as long as e and e' are not switched, we would have e and e' both at position (s, t_2) . (There might be some -1 before D_i 's too, but this can be neglected.) The example given after this proof illustrates this calculation.

We then apply other $\pi_{i,j}$'s to $C^{\pi_{1,1}}$ sequentially; each application of $\pi_{i,j}$ would yield a bunch of D_i 's in $\mathfrak{M}(n, d)$. At last we shall get $C = \text{sgn}(\sigma)C^\sigma + \sum_i D_i$ where $D_i \in \mathfrak{M}(n, d)$. For every $\sigma \in S_T$ such an equation can be derived (when $\sigma = \text{id}$ use $C = C$), and we have $|T|!C = \sum_{\sigma \in S_T} \text{sgn}(\sigma)C^\sigma + \sum_i D_i$. Plugging the second type of relations in, and noting that $|T| > n^2$, $C = \frac{1}{|T|!}(\sum_i D_i)$ where $D_i \in \mathfrak{M}$ completing the proof. \square

Example 13. In this example

$$S = \begin{vmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \end{vmatrix}, \hat{T} = \begin{vmatrix} \hat{1} & \hat{3} \\ \hat{2} & \hat{5} \\ \hat{4} & \hat{6} \end{vmatrix}.$$

The corresponding graph is



The kernel has the following relation coming from the right monomial.

$$\begin{vmatrix} \hat{1} & \hat{3} \\ \hat{2} & \hat{5} \\ \hat{4} & \hat{6} \end{vmatrix} = \begin{vmatrix} \hat{2} & \hat{3} \\ \hat{1} & \hat{5} \\ \hat{4} & \hat{6} \end{vmatrix} + \begin{vmatrix} \hat{1} & \hat{2} \\ \hat{3} & \hat{5} \\ \hat{4} & \hat{6} \end{vmatrix}$$

Multiplying this by S and recalling our convention on associating monomials to correlated tableaux, we have

To illustrate the utility of Proposition 12 developed so far, we consider matrix semi-invariants for 2×2 matrices. For 2×2 matrices, note that as soon as there exists a multiple edge, a 2-regular bipartite graph is disconnected. Therefore, an immediate application of Proposition 12 gives the following known result.

Theorem 14. *The matrix semi-invariants of 2×2 matrices are generated by those of degree ≤ 4 .*

For 2×2 matrices, Domokos presented an explicit generating set, and from this description he deduced that $\beta = 4$ [Dom00a]. Therefore our bound is tight in this case. Also note that our result is obtained without computing a single invariant.

Acknowledgement. We would like to thank Mátyás Domokos, Bharat Adsul, Ketan Mulmuley, Partha Mukhopadhyay and K. N. Raghavan for discussions related to this work. Part of the work was done when Youming was visiting the Simons Institute for the program Algorithms and Complexity in Algebraic Geometry.

References

- [ANS07] B. Adsul, S. Nayak, and K. V. Subrahmanyam. A geometric approach to the Kronecker problem II: rectangular shapes, invariants of matrices and the Artin–Procesi theorem. preprint, 2007.
- [BD06] M. Bürgin and J. Draisma. The Hilbert null-cone on tuples of matrices and bilinear forms. *Mathematische Zeitschrift*, 254(4):785–809, 2006.

- [CIW97] Ajeh M Cohen, Gábor Ivanyos, and David B Wales. Finding the radical of an algebra of linear transformations. *Journal of Pure and Applied Algebra*, 117:177–193, 1997.
- [DD12] M. Domokos and V. Drensky. Defining relation for semi-invariants of three by three matrix triples. *Journal of Pure and Applied Algebra*, 216(10):2098–2105, 2012.
- [Der01] Harm Derksen. Polynomial bounds for rings of invariants. *Proceedings of the American Mathematical Society*, 129(4):955–964, 2001.
- [DKZ02] M. Domokos, S. G. Kuzmin, and A. N. Zubkov. Rings of matrix invariants in positive characteristic. *Journal of Pure and Applied Algebra*, 176(1):61–80, 2002.
- [Dom00a] M. Domokos. Poincaré series of semi-invariants of 2×2 matrices. *Linear Algebra and its Applications*, 310(1):183–194, 2000.
- [Dom00b] M. Domokos. Relative invariants of 3×3 matrix triples. *Linear and Multilinear Algebra*, 47(2):175–190, 2000.
- [Dom02] M. Domokos. Finite generating system of matrix invariants. *Math. Pannon*, 13(2):175–181, 2002.
- [Dom03] M. Domokos. *Matrix invariants and the failure of Weyl’s theorem*, volume 235 of *Lecture Notes in Pure and Applied Mathematics*, page 215–236. Dekker, New York, 2003.
- [Don92] Stephen Donkin. Invariants of several matrices. *Inventiones mathematicae*, 110(1):389–401, 1992.
- [Don93] Stephen Donkin. Invariant functions on matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 113:23–43, 1 1993.
- [DW00] Harm Derksen and Jerzy Weyman. Semi-invariants of quivers and saturation for littlewood-richardson coefficients. *Journal of the American Mathematical Society*, 13(3):467–479, 2000.
- [DZ01] M. Domokos and A. N. Zubkov. Semi-invariants of quivers as determinants. *Transformation groups*, 6(1):9–24, 2001.
- [For86] Edward Formanek. Generating the ring of matrix invariants. In Freddy M. J. van Oystaeyen, editor, *Ring Theory*, volume 1197 of *Lecture Notes in Mathematics*, pages 73–82. Springer Berlin Heidelberg, 1986.
- [Ful97] W. Fulton. *Young Tableaux: With Applications to Representation Theory and Geometry*. London Mathematical Society Student Texts. Cambridge University Press, 1997.
- [Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. Syst. Sci.*, 69(3):448–484, 2004.
- [HW14] Pavel Hrubeš and Avi Wigderson. Non-commutative arithmetic circuits with division. In *Innovations in Theoretical Computer Science, ITCS’14, Princeton, NJ, USA, January 12-14, 2014*, pages 49–66, 2014. Full version at <http://www.math.ias.edu/~avi/PUBLICATIONS/HrubesWiDec2013.pdf>.

- [IQS15] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Non-commutative Edmonds' problem and matrix semi-invariants. *CoRR*, abs/1508.00690, 2015.
- [Pro76] C. Procesi. The invariant theory of $n \times n$ matrices. *Advances in Mathematics*, 19(3):306–381, 1976.
- [Raz74] Ju. P. Razmyslov. Trace identities of full matrix algebras over a field of characteristic zero. *Mathematics of the USSR-Izvestiya*, 8(4):727, 1974. English translation available at <http://iopscience.iop.org/0025-5726/8/4/A01>.
- [Sag01] B.E. Sagan. *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*. Graduate Texts in Mathematics. Springer, 2001.
- [SVdB01] Aidan Schofield and Michel Van den Bergh. Semi-invariants of quivers for arbitrary dimension vectors. *Indagationes Mathematicae*, 12(1):125–138, 2001.
- [Wey97] H. Weyl. *The classical groups: their invariants and representations*, volume 1. Princeton University Press, 1997.